

CLAIMS

What is claimed is:

- 1 1. A method comprising:
 - 2 locating blocks of data in a log that are referenced and within a range at a tail of
 - 3 the log; and
 - 4 copying the blocks of data that are referenced and within the range to an
 - 5 unallocated segment of the log.
- 1 2. The method of claim 1 comprising designating the range as unallocated.
- 1 3. The method of claim 1, wherein the blocks of data are associated with nodes in a storage tree within the log and wherein locating the blocks of data that are referenced and within the range includes determining a minimum value among addresses of descendent nodes of a node.
- 1 4. The method of claim 3, wherein a location table includes an entry for nodes that reference other nodes and wherein determining the minimum value among addresses of descendent nodes of the node include retrieving the minimum value from an entry in the location table associated with the node.
- 1 5. The method of claim 4, wherein locating the blocks of data that are referenced and within the range includes processing the descendent nodes of the node upon determining that the minimum value among the addresses of the descendent nodes is within the range.
- 1 6. The method of claim 5 comprising modifying the addresses of the copied blocks of data that are stored in the location table based on the new locations of the copied blocks of data in the log.

1 7. The method of claim 5 comprising modifying the minimum value in the entry in
2 the table associated with the node when the minimum value changes based the new
3 locations of the copied blocks of data that are associated with descendent nodes of the
4 node.

1

1 8. A method comprising:

2 garbage collecting within a range of addresses in a storage system, which includes
3 a plurality of storage trees having multiple references to the same block of
4 data, by

5 pruning walking of the plurality of storage trees to determine active blocks
6 of data within said range, where active blocks of data are those still
7 in one of the plurality of storage trees, by

8 determining, based on accessing in one of said plurality of storage
9 trees a parent node that has a plurality of descendent nodes,
10 that none of the plurality of descendant nodes are associated
11 with blocks of data within the range; and

12 skipping the walking of the plurality of descendent nodes based on
13 said determining.

1

1 9. The method of claim 8, wherein the blocks of data are stored in a log and the
2 range is a segment of the log.

1

1

1 10. The method of claim 9, wherein the segment is at the tail of the log.

1

1 11. The method of claim 10, wherein the determining is performed by comparing a
2 minimum offset of the plurality of descendent nodes against the range, wherein the

3 minimum offset is accessed when walking the parent node and without walking the
4 plurality of descendent nodes.

1

1 12. The method of claim 8, wherein the garbage collecting is further performed by:
2 copying the active blocks of data out of the range; and
3 marking the range as unallocated.

1

1 13. The method of claim 12, wherein the range is a segment at the tail of a log and
2 said copying is from the said segment at the tail to a segment at the head of the log.

1

1 14. The method of claim 12, wherein said copying includes updating addresses of the
2 copied blocks of data within a location table.

1 15. A method comprising:

2 performing the following until each block of data that is active in a range to be
3 cleaned at a tail of a log of data is copied to a head of the log, wherein a block of data is
4 associated with a node of a storage tree,

5 copying blocks of data associated with child nodes of a current node that
6 are within the range to be cleaned to the head of the log;

7 retrieving a block of data associated with the current node, upon
8 determining that a minimum address value among addresses of descendent nodes is
9 within the range to be cleaned;

10 designating, as the current node, one of the child nodes of the current node
11 that is an interior node, upon determining that at least one child node is an interior node;
12 and

13 designating, as the current node, an ancestor node of the current node
14 whose descendent nodes are unprocessed.

1 16. The method of claim 15, wherein performing the following until each block of
2 data that is active in the range to be cleaned at the tail of the log of data is copied to a
3 head of the log includes updating addresses of that copied blocks of data within a location
4 table.

1 17. The method of claim 15, wherein performing the following until each block of
2 data that is active in the range to be cleaned at the tail of the log of data is copied to the
3 head of the log includes updating a minimum address value among addresses of
4 descendent nodes for an entry for the current node in a location table when the minimum
5 address value changes based on copying of the blocks of data associated with the
6 descendent nodes of the current node.

1 18. The method of claim 15, wherein at least one block of data stored in the log is
2 referenced by more than one of other blocks of data.

1 19. The method of claim 15 comprising marking the range as unallocated when the
2 blocks of data that are active and within the range are copied to the head of the log.

1 20. A system comprising:
2 a storage device to store a number of blocks of data, wherein the blocks of data
3 that are marked as allocated are non-modifiable, the blocks of data to be stored as a log;
4 and
5 a garbage collection logic to locate the blocks of data that are referenced and
6 within a range at a tail of the log.

1 21. The system of claim 20, wherein the garbage collection logic is to copy the blocks
2 of data that are referenced to an unallocated address space of the log.

1 22. The system of claim 21, wherein the garbage collection logic is to copy the blocks
2 of data that are referenced to a head of the log.

1 23. The system of claim 20, wherein the garbage collection logic is to mark the range
2 as unallocated.

1 24. The system of claim 20, wherein at least one of the number of blocks of data are
2 referenced by more than one reference.

1 25. The system of claim 20 comprising a location table to include entries associated
2 with interior nodes of a storage tree, wherein each entry is to include a minimum value
3 among the addresses of descendent nodes of the associated interior node.

1 26. The system of claim 25, wherein the garbage collection logic is to locate the
2 blocks of data that are referenced and within the range at the tail of the log based on the
3 minimum values stored in the entries of the location table.

1 27. A backup system comprising:
2 a set of one or more storage trees, each representing a snapshot of a file system at
3 a different time, each leaf node of said set of storage trees to include a block of data from
4 said file system that has been backed up from a set of one or more storage devices;
5 a storage space to store said blocks of data having been allocated from a backup
6 storage space in said set of storage devices;
7 a set of one or more location tables having stored therein a minimum address
8 value for descendent nodes of interior nodes of said set of storage trees; and
9 a garbage collection logic to clean a currently selected range from the tail of said
10 log , said garbage collection logic to prune walking of nodes of said set of storage trees
11 based on said set of location tables and said currently selected range.

1 28. The backup system of claim 27, wherein two different nodes of a same storage
2 tree reference a same node in the same storage tree.

1 29. The backup system of claim 27, wherein the garbage collection logic is to update
2 references to a node that is within the currently selected range based on an update to an
3 entry in the set of one or more location tables.

1 30. The backup system of claim 27, wherein the garbage collection logic is to prune
2 walking of the nodes of said set of storage trees based on the minimum addresses stored
3 in the set of one or more location tables.

1 31. An apparatus comprising:
2 a backup system to backup a file system, said backup file system including,
3 a tracking logic to generate a set of trees each representing backup
4 snapshots of said file system at different times by recording references to blocks of
5 backed up data stored in a set of one more storage devices;
6 an allocator logic to allocate contiguous blocks of storage space from a log
7 of a back up storage space to store said blocks of backed up data;
8 a garbage collection logic to, responsive to deletion of one or more of said
9 backup snapshots, clean a currently selected contiguous range from the tail of said log,
10 said garbage collection logic to,
11 walk only those nodes of said set of storage trees that possibly
12 identify those of said blocks of data that are stored in said currently
13 selected contiguous range or that possibly are themselves stored in said
14 currently selected contiguous range, and
15 sweep said currently selected contiguous range.

1 32. The apparatus of claim 31, wherein the set of trees include interior nodes and leaf
2 nodes, the interior nodes to include references to other nodes in the set of one or more
3 storage trees, two different interior nodes of a same tree references a same node in the
4 same tree.

1 33. A machine-readable medium that provides instructions, which when executed by a
2 machine, cause said machine to perform operations comprising:

3 locating blocks of data in a log that are referenced and within a range at a tail of
4 the log; and

5 copying the blocks of data that are referenced and within the to an unallocated
6 segment of the log.

1 34. The machine-readable medium of claim 33 comprising designating the range as
2 unallocated.

1 35. The machine-readable medium of claim 33, wherein the blocks of data are
2 associated with nodes in a storage tree within the log and wherein locating the blocks of
3 data that are referenced and within the range includes determining a minimum value
4 among addresses of descendent nodes of a node.

1 36. The machine-readable medium of claim 35, wherein a location table includes an
2 entry for nodes that reference other nodes and wherein determining the minimum value
3 among addresses of descendent nodes of the node include retrieving the minimum value
4 from an entry in the location table associated with the node.

1 37. The machine-readable medium of claim 36, wherein locating the blocks of data
2 that are referenced and within the range includes processing the descendent nodes of the

3 node upon determining that the minimum value among the addresses of the descendent
4 nodes is within the range.

1 38. The machine-readable medium of claim 37 comprising modifying the addresses of
2 the copied blocks of data that are stored in the location table based on the new locations
3 of the copied blocks of data in the log.

1 39. The machine-readable medium of claim 37 comprising modifying the minimum
2 value in the entry in the table associated with the node when the minimum value changes
3 based the new locations of the copied blocks of data that are associated with descendent
4 nodes of the node.

1

1 40. A machine-readable medium that provides instructions, which when executed by a
2 machine, cause said machine to perform operations comprising:

3 garbage collecting within a range of addresses in a storage system, which includes
4 a plurality of storage trees having multiple references to the same block of
5 data, by

6 pruning walking of the plurality of storage trees to determine active blocks
7 of data within said range, where active blocks of data are those still
8 in one of the plurality of storage trees, by

9 determining, based on accessing in one of said plurality of storage
10 trees a parent node that has a plurality of descendent nodes,
11 that none of the plurality of descendent nodes are associated
12 with blocks of data within the range; and

13 skipping the walking of the plurality of descendent nodes based on
14 said determining.

1

1 41. The machine-readable medium of claim 40, wherein the blocks of data are stored
2 in a log and the range is a segment of the log.

1 42. The machine-readable medium of claim 41, wherein the segment is at the tail of
2 the log.

1

1 43. The machine-readable medium of claim 42, wherein the determining is performed
2 by comparing a minimum offset of the plurality of descendent nodes against the range,
3 wherein the minimum offset is accessed when walking the parent node and without
4 walking the plurality of descendent nodes.

1

1 44. The machine-readable medium of claim 40, wherein the garbage collecting is
2 further performed by:

3 copying the active blocks of data out of the range; and
4 marking the range as unallocated.

1

1 45. The machine-readable medium of claim 44, wherein the range is a segment at the
2 tail of a log and said copying is from the said segment at the tail to a segment at the head
3 of the log.

1

1 46. The machine-readable medium of claim 44, wherein said copying includes
2 updating addresses of the copied blocks of data within a location table.

1 47. A machine-readable medium that provides instructions, which when executed by a
2 machine, cause said machine to perform operations comprising:

3 performing the following until each block of data that is active in a range to be
4 cleaned at a tail of a log of data is copied to a head of the log, wherein a block of data is
5 associated with a node of a storage tree,

6 copying blocks of data associated with child nodes of a current node that
7 are within the range to be cleaned to the head of the log;

8 retrieving a block of data associated with the current node, upon
9 determining that a minimum address value among addresses of descendent nodes is
10 within the range to be cleaned;
11 designating, as the current node, one of the child nodes of the current node
12 that is an interior node, upon determining that at least one child node is an interior node;
13 and
14 designating, as the current node, an ancestor node of the current node
15 whose descendent nodes are unprocessed.

1 48. The machine-readable medium of claim 47, wherein performing the following
2 until each block of data that is active in the range to be cleaned at the tail of the log of
3 data is copied to a head of the contiguous log includes updating addresses of that copied
4 blocks of data within a location table.

1 49. The machine-readable medium of claim 47, wherein performing the following
2 until each block of data that is active in the segment to be cleaned at the tail of a log of
3 data is copied to the head of the log includes updating a minimum address value among
4 addresses of descendent nodes for an entry for the current node in a location table when
5 the minimum address value changes based on copying of the blocks of data associated
6 with the descendent nodes of the current node.

1 50. The machine-readable medium of claim 47, wherein at least one block of data
2 stored in the log is referenced by more than one of other blocks of data.

1 51. The machine-readable medium of claim 47 comprising marking the range as
2 unallocated when the blocks of data that are active and within the range are copied to the
3 head of the log.